Distributed Display Environments

Dugald Ralph Hutchings, John Stasko
College of Computing/GVU Center
Georgia Institute of Technology
Atlanta, GA 30332
{hutch, stasko}@cc.gatech.edu

Mary Czerwinski
Microsoft Research
One Microsoft Way
Redmond, WA 98052 USA
marycz@microsoft.com

INTRODUCTION

The term "Distributed Display Environment" (DDE) compactly describes computer systems that present output to more than one physical display. Often DDEs allow attached input devices to treat the displays as a contiguous surface but this is not necessarily a requirement. In much previous HCI research, there are implicit or explicit assumptions that a computer system has at most one physical display attached. However, very recent research has recognized that a computer system can have more than one physical display. Research methods and foci have varied, but findings suggest that (1) continuing to define the relevant aspects of DDEs is worthwhile due to obvious and quantifiable benefits of these environments and (2) there are important aspects of interfaces that have been ignored in the past because of the assumption of a single-display output model.

DDE RESEARCH: PAST AND PRESENT

In this proposal, we will discuss some past work in DDEs. The intent of the analysis of past work is to motivate the need and desirability for a workshop on this topic. Following the analysis we will describe a host of questions surrounding DDEs.

The earliest research in the computer science literature regarding multiple-monitor systems actually occurred about ten years ago when intelligent systems researchers Funke, Neal, and Paul developed an intelligent window manager for a dual-monitor system [3]. They attempted to automate the showing and hiding of windows by measuring, among other properties, relative position of windows and amount of time that windows were active. A simple evaluation of the system revealed that windows were often incorrectly shown or hidden and, in particular, windows being used as reference material for interaction in a focus window were often unexpectedly hidden. At a higher level, this piece of research provides an example of difficulties that can arise when assumptions about computer use of single-display systems carry over into systems designed to be used in DDEs.

Much of the recent research related to DDEs touches on this topic of *information display* rather than interaction. Grudin exposes several classes of use for multiple-monitor desktop computer systems including a class of users who dedicate secondary space to the display of communication agents [4]. Ringel followed this work by comparing the use of virtual desktops against Grudin's documented uses of multiple monitors, finding a class of users who *avoid* multiple-monitor systems because the communication information that is displayed on secondary monitors distracts them [10]. Hutchings *et al*. confirmed the analysis given by Grudin with a quantitative demonstration that multiple-monitor users often had their email windows visible but inactive, *i.e*. not receiving input from a user [6]. Mackinlay and Heer provide a software system that allows users to overcome the physical separation of displays in a DDE to perceive text and images as if they were being displayed on a single, large-pixel display, thus dedicating the entire space to information display [8], though further studies have shown a wider variety of multiple-monitor information layout and usage patterns [6]. MacIntyre *et al*. and Robertson *et al*. explicitly segment the DDE into a focus region, where the current task is being conducted, and an awareness region, where information from other tasks is shown [7, 11]. Perry and O'Hara show how people use their desks and walls to arrange objects for information display and suggest that these observations can be applied to the design of interfaces for DDEs [9].

Many of the aforementioned pieces of research also address interaction: Grudin shows a class of users who place UI tools and widgets on one monitor while placing information on another [4]; Ringel's findings [10] suggest that virtual desktops could be revised to address concerns of distracting information; Hutchings *et al*. showed that interaction components like the MS-Windows TaskBar tend to get much less use by multiple-monitor users [6]. Notice though that all of these studies still relate to the *placement* or *location* of on-screen objects and information. These observations have led us to two conclusions: (1) understanding and designing for the display of information in non-active display regions is considerably more important than in previous times and (2) there is a need to recognize opportunities for improvement and innovation in *interaction* and move beyond research that focuses only on information display for DDEs.

Of course the discussion about aspects of DDEs is irrelevant unless people intrigued by the concept can demonstrate that DDEs are worthwhile to users. Initial evidence about multiple monitor systems indicates that (1) even without alterations to basic interaction, users can complete tasks more quickly with multiple monitors, and (2) users typically experience inefficiencies or frustrations that can be addressed by alterations to basic interaction. Czerwinski *et al*. and Tan & Czerwinski conducted studies that provide this evidence [2, 13]. Private research also provides some evidence, some of which was conducted for companies that benefit from such evidence[1] and others that had no stated interest in the results.[2, 3]

WORKSHOP GOALS

Our analysis of the recent work in the area of DDEs shows two distinct areas of consideration for future research: design for information display and design for interaction. In addition to developing goals around these two aspects, we feel that we can strengthen the community of people conducting research in this area by addressing evaluation so that we have a common understanding of how to assess technological innovations and other research activities in the area. Finally, we think it is worthwhile to discuss some of the boundaries of work in DDEs as well as some of the broader implications that our findings may be able to bring to HCI research in general.

1. **Determine promising areas for new *interaction research* about DDEs**. While much of the recent research involving DDEs has focused on information display, very little has examined interaction itself. There are three main subgoals of this goal: (a) examine areas of basic interaction (such as cursor movement) in need of revision or innovation for DDEs, (b) identify those research that are most likely to be strongly affected by DDEs (such as information visualization, perhaps), and (c) discuss whether, and to what degree, well-known interaction techniques can be applied to issues of interaction on DDEs (for example, Magic Lenses [1] are not common features of many interfaces, but would they improve interaction in DDEs? If so, how strong is the impact likely to be?).

2. **Detail the aspects of incorporating *passive information display* into interfaces**. The recent research indicates that people seem to dedicate more space to the display of often-used (in the case of email) or highly task-relevant (in the case of related web-based material) information. As a group, it will be helpful to discuss to what degree this trend will continue and whether the size and orientation of displays in DDEs allows us to predict how much space should be devoted simply to displayed information. Furthermore, we can discuss areas of interface development related to this trend, such as automatic alteration of displayed contents to optimally display information, classifying both the type of displayed information and the appropriate techniques for handling each class, and considering interfaces that "more calmly" display information (such as Stasko's *et al*. InfoCanvas [12]). In particular, a discussion of what computer technology capabilities and interfaces can provide beyond static information display could suggest new areas of interface development.

3. **Evaluation: determining the "what" and the "how."** The community could benefit from a discussion of what to evaluate. Recent evaluation work has looked primarily at comparing single-monitor to multiple-monitor use for some small set of tasks. Debates could ask the question "Should we focus on more tasks and naïve comparisons to determine task areas that need attention, or should we target likely areas right away and compare multiple-monitor users with and without the newer interfaces?" There is also a question of generalizability: Does a comparison of a single-monitor scenario and a dual-monitor scenario apply to three-monitors or four-monitors? Does the orientation and physical placement of monitors make a difference? How do differently-sized monitors in a DDE differ from a DDE in which every monitor is the same size? There is also the question of "how"; are new evaluation techniques required in order for testing interfaces and display configurations, and if so, how do we ascertain how confidently we can take the results from the new methods?

4. **Outline broader implications to HCI research.** Initial findings in the area of DDEs raise broader questions. For example, does direct manipulation need major reconsideration when there are a huge number of pixels to traverse? How does Fitts' Law apply to a space that is physically broad but virtually narrow? How much of our understanding of visual perception as it pertains to HCI depends on the existence of a single monitor? Research in DDEs provides the opportunity to challenge common models of interaction and this workshop should strive to (1) identify those models that are the most likely to be susceptible to "one-monitor bias" and (2) suggest experiments or other research that can determine how robust the models actually are.

DETAILED PLAN FOR CONDUCTING THE WORKSHOP

As suggested by the proposal submission website for CHI workshops, there are three parts to the detailed plan for conducting the workshop: activities to be conducted prior to the workshop, structures for the workshop itself, and follow-up.

Before the Workshop

We would like the participants to have read each other's workshop position statements and a few other papers we will put on the website as background material. There will be time to discuss and review these findings during the first 90 minutes of the workshop. Furthermore, we will encourage participants to suggest aspects of the four goals that they wish to discuss prior to arriving at the workshop. This will help us to prepare materials to guide the discussions of the general topics.

Schedule of the Workshop

We begin this section with a projected timetable for the workshop.

09:00	Opening comments, introductions, etc.	13:00	Lunch
09:15	Brief presentations and discussion	14:00	Group discussion on goal 3
10:30	Coffee break	15:00	Group discussion on goal 4
11:00	Group discussion on goal 1	16:00	Coffee break
12:00	Group discussion on goal 2	16:30	Wrap-up, conclusions, closing comments

We are aiming for a 6-hour, 1 day workshop in which we devote the majority of the time to addressing the stated goals. With a small group of approximately 12 participants, we should be able to make progress as an entire group in attacking each of the four goals in the designated one-hour periods. The schedule yields the flexibility, however, to split into smaller groups to attack subgoals of the main goals if we find interest in particular subgoals among a subgroup of participants or if the general discussion is too unwieldy. In the final thirty minutes, we plan to gather consensus on the most important discussions, findings, and conclusions.

After the Workshop

Our plan for dissemination of the workshop results will be to publish the progress we make in the SIGCHI Bulletin, in addition to updating the website with discussion results and making the URL widely available to the HCI community. In addition, if there is interest among attendees, it might make sense to explore a special journal issue based on key research directives identified or pursued by attendees.

ORGANIZERS BACKGROUNDS

Duke Hutchings is a Ph.D. Candidate in John Stasko's Information Interfaces research group and has held two internships with Mary Czerwinski's Visualization and Interaction research group at Microsoft Research. His work with those groups focuses on window management issues for multiple-monitor systems [5, 6, 11]. More information is available on his website: http://www.cc.gatech.edu/~hutch.

Dr. John Stasko is a Professor in the College of Computing and the GVU Center at the Georgia Institute of Technology, where he is Director of the Information Interfaces Research Group (http://www.cc.gatech.edu/gvu/ii). He received his Ph.D. degree in computer science from Brown University in 1989 and joined the faculty at Georgia Tech that same year. His research is in the area of Human-Computer Interaction with a specific focus on information visualization, peripheral awareness of information, and software agents. He has been author or co-author of over 25 journal articles and 60 conference papers, PI or co-PI on one ONR and seven NSF research grants. He is on the editorial staff of the International Journal of Human-Computer Studies, Journal of Visual Languages and Computing, Information Visualization, and World Wide Web. He was Program Chair for the 2003 ACM Symposium on Software Visualization and Program Co-Chair for the 2000 IEEE Symposium on Visual Languages.

Dr. Mary Czerwinski is a Senior Researcher and Manager of the Visualization and Interaction Research group at Microsoft Research. The group is responsible for studying and designing advanced technology and interaction techniques that leverage human capabilities across a wide variety of input and output channels. Mary's primary research areas include spatial cognition, information visualization and task switching. Mary has been an affiliate assistant professor at the Department of Psychology, University of Washington since 1996. She has also held positions at Compaq Computer Corporation, Rice University, Lockheed Engineering and Sciences Corporation, and Bell Communications Research. She received a Ph.D. in cognitive psychology from Indiana University in Bloomington. Mary is active in the field of Human-Computer Interaction, publishing and participating in a wide number of conferences, professional venues and journals. More information about Dr. Czerwinski can be found on her webpage: http://research.microsoft.com/users/marycz.

REFERENCES

1. Bier, E. A. et al. Toolglass and magic lenses: the see-through interface. Proc. CHI 1993, ACM Press, 73 80.
2. Czerwinski, M. et al. Toward characterizing the productivity benefits of very large displays. Proc. Interact 2003, IOS Press, 9 16.
3. Funke, D. Neal, J., and Paul, R. An approach to intelligent automated window management. Int. J. of Man-Machine Studies 38, 1993, 949 983.
4. Grudin, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In Proc. CHI 2001, ACM Press, 458 465.
5. Hutchings, D. R., and Stasko, J. Revisiting Display Space Management: Understanding Current Practice to Inform Next-generation Design. Proc. Graphics Interface 2004. Canadian Human-Computer Communications Society, 127 134.
6. Hutchings, D. R., et al. Display space usage and window management operation comparisons between single monitor and multiple monitor users. Proc. Advanced Visual Interfaces 2004, ACM Press, 32 39.
7. MacIntyre, B., et al. Support for multitasking and background awareness using interactive peripheral displays. Proc. UIST 2001, ACM Press, 41 50.
8. Mackinlay, J., and Heer, J. Wideband displays: mitigating multiple monitor seems. CHI Extended Abstracts 2004, ACM Press, 1521 1524.
9. Perry, M. and O'Hara, K. Display-based Activity in the workplace. Proc. Interact 2003, IOS Press, 591 598.
10. Ringel, M. When one isn't enough: an analysis of virtual desktop usage strategies and their implications for design. CHI Extended Abstracts 2003, ACM Press, 762 763.
11. Robertson, G. et al. Scalable Fabric: Flexible task management. Proc. Advanced Visual Interfaces 2004, ACM Press, 85 89.
12. Stasko, J. et al. Personalized peripheral information awareness through information art. Proc. UbiComp 2004, Springer-Verlag, 18 35.
13. Tan, D.S. and Czerwinski, M. Effects of visual separation and physical continuities when distributing information across multiple displays. Proc. Interact 2003, IOS Press, 252 265.

[1] ATI and NEC commissioned research to this end: http://www.eetimes.com/showPressRelease.jhtml?articleID=103399
[2] A casual experiment: http://dotnetjunkies.com/WebLog/darrell.norton/archive/2003/11/11/3432.aspx
[3] A more formal set of analyses: "Multiple Display Market and Consumer Attitudes" from Jon Peddie Research. Main site: http://jonpeddie.com